Chapter 2
Basics of Game Theory

Abstract This chapter provides a brief overview of basic concepts in game theory. These include game formulations and classifications, games in extensive vs. in normal form, games with continuous action (strategy) sets vs. finite strategy sets, mixed vs. pure strategies, and games with uncoupled (orthogonal) vs. coupled action sets. The next section reviews basic solution concepts, among them Nash equilibria being of most relevance. The chapter is concluded with some remarks on the rationality assumption and learning in classical games. The following chapters will introduce these concepts formally.

2.1 Introduction

Game theory is a branch of applied mathematics concerned with the study of situations involving conflicting interests. Specifically game theory aims to mathematically capture behavior in strategic situations, in which an individual's success in making choices depends on the choices of others. The field was born with the work of John von Neumann and Oskar Morgenstern [159], although the theory was developed extensively in the 1950s by many among whom John Nash [95, 96].

In this chapter we shall introduce the game-theoretic notions in simplest terms. Our goal will be later on to study and formalize mathematically various game problems, by which we understand problems of conflict with common strategic features. While initially developed to analyze competitions in which one individual does better at another's expense (zero-sum games), it has been expanded to treat a wide class of interactions, which are classified according to several criteria, one of these being cooperative versus noncooperative interactions. Typical classical games are used to model and predict the outcome of a wide variety of scenarios involving a finite number of players (or agents) that seek to optimize some individual objective. Noncooperative game theory studies the strategic interaction among self-interested players.

Historically, game theory developments were motivated by studies in economics, but many interesting game theory applications have emerged in fields as diverse as biology [141], computer science [54], social science and engineering [74]. In engineering, the interest in noncooperative game theory is motivated by the possibility of designing large scale systems that globally regulate their performance in a dis-

L. Pavel, *Game Theory for Control of Optical Networks*,
Static & Dynamic Game Theory: Foundations & Applications,
DOI 10.1007/978-0-8176-8322-1_2,

tributed, and decentralized manner. Modeling a problem within a game-theoretic setting is particularly relevant to any practical application consisting of separate subsystems that compete for the use of some limited resource. Examples of such applications include most notably congestion control in network traffic (i.e. Internet, or transportation), problems of optimal routing [11, 13, 14], power allocation in wireless communications and optical networks [118, 133].

The chapter is organized as follows. We present basic concepts and game formulations, then we review some classifications: games in extensive vs. in normal form, games with continuous (infinite) action (strategy) set vs. finite action (finite strategy games), mixed vs. pure strategy games, and games with uncoupled (orthogonal) vs. coupled action sets. We follow with a discussion of basic solution concepts, among them Nash equilibria being of most relevance herein. We conclude with some remarks on the rationality assumption and learning in classical games.

2.2 Game Formulations

Game theory involves multiple decision-makers and sees participators as competitors (players). In each game players have a sequence of personal moves; at each move, each player has a number of choices from among several possibilities, also possible is the chance or random move. At the end of the game there is some payoff to be gained (cost to be paid) by the players which depends on how the game was played. Noncooperative game theory [96] studies the strategic interaction among self-interested players. A game is called *noncooperative* if each player pursues its own interests which are partly conflicting with others'. It is assumed that each player acts independently without collaboration or communication with any of the others [96].

This is in contrast to a standard optimization where there is only one decision-maker who aims to minimize an objective function by choosing values of variables from a constrained set such that the system performance is optimized (see the appendix for a review of basic results, mainly drawn from [24]).

So far we have mentioned three elements: alternation of moves (individual or random (chance)), a possible lack of knowledge and a payoff or cost function. A game $\mathcal{G}$ consists of a set of players (agents) $\mathcal{M} = \{1, \dots, m\}$, an action set denoted Ω_i (also referred to as a set of strategies S_i) available for each player i and an individual payoff (utility) $\mathcal{U}_i$ or cost function J_i for each player $i \in \mathcal{M}$.

In a game, each player individually takes an optimal action which optimizes its own objective function and each player's success in making decisions depends on the decisions of the others. We define a noncooperative game $\mathcal{G}$ as an object specified by the triplet $(\mathcal{M}, S, \Omega, \mathbf{J})$, where

$$S = S_1 \times S_2 \times \cdots \times S_m$$

is known as the strategy space,

$$\Omega = \Omega_1 \times \Omega_2 \times \cdots \times \Omega_m$$

is the action space, and $\mathbf{J} : \Omega \to R^m$, defined as

$$\mathbf{J}(\mathbf{u}) = \left[J_1(\mathbf{u}), \ldots, J_m(\mathbf{u})\right]^T, \quad \mathbf{u} \in \Omega$$

is the vector of objective functions associated to each of the m players, or agents participating in the game. In some cases (see Chap. 12) a graph notation might be more appropriate than the set $\mathcal{M}$ notation. Conventionally $\mathbf{J}$ represents a vector of cost functions to be minimized by the agents. An alternative formulation is to use a vector of utility functions (payoff functions) $\mathcal{U}$ to be maximized by the agents. Without loss of generality we assume that each agent aims to minimize its cost so we will use $\mathbf{J}$ throughout. Since usually we shall identify each player's action set and cost function we shall use the notation $\mathcal{G}(\mathcal{M}, S_i, \Omega_i, J_i)$, where the subscript is associated to each player i, $i \in \mathcal{M}$. In some cases we shall identify strategies with actions (one-shot games) and drop one of the arguments. The properties of the sets Ω_i and the functions $J_i(\mathbf{u})$, $i \in \mathcal{M}$ depend on the modeling scenario, and hence the type of game under consideration.

Each player's success in making decisions depends on the decisions of the others. Let Ω_i denote the set of actions available to player i, which can be finite or infinite. This leads to either finite actions set games, also known as matrix games, or infinite (continuous action set) games. In the latter case each player can choose its action from a continuum of (possibly vector-valued) alternatives. A strategy can be regarded as a rule for choosing an action, depending on external conditions. Once such a condition is observed, the strategy is implemented as an action. In the case of *mixed strategies*, this external condition is the result of some randomization process. Briefly, a *mixed strategy* for agent i is a probability distribution x^i over its action set Ω_i. In some cases actions are *pure*, or independent of any external conditions, and the strategy space coincides with the action space. In discussing games in *pure strategies* we shall use the term "strategy" and "action" interchangeably to refer to some $\mathbf{u} \in \Omega$, and the game $\mathcal{G}$ can simply be specified as the pair $\mathcal{G}(\mathcal{M}, \Omega_i, J_i)$.

In the introduction we have already distinguished between cooperative and noncooperative games. There are numerous other classifications of games, but only a few are relevant to our purposes in this monograph. We will briefly review the distinctions between these types of game and introduce these concepts for possible forms of a game as well as what we understand by various solution concepts.

2.3 Games in Extensive Form

The extensive form of a game amounts to a translation of all the rules into technical terms of a formal system designed to describe all games.

Extensive-form games generally involve several acts or stages, and each player chooses a strategy at each stage. The game's information structure, i.e., how much information is revealed to which players concerning the game's outcomes and their opponents' actions in the previous stages, significantly affects the analysis of such games. Extensive-form games are generally represented using a tree graph. Each node (called a decision node) represents every possible state of play of the game as it is played [20]. Play begins at a unique initial node, and flows through the tree along a path determined by the players until a terminal node is reached, where play ends and cost s are assigned to all players. Each non-terminal node belongs to a player; that player chooses among the possible moves at that node, each possible move is an edge leading from that node to another node. Their analysis becomes difficult with increasing numbers of players and game stages.

A formal definition is as follows.

Definition 2.1 An *m-player game $\mathcal{G}$ in an extensive form* is defined as a graph theoretic tree of vertices (states) connected by edges (decisions or choices) with certain properties:

1. $\mathcal{G}$ has a specific vertex called the *starting point* of the game,
2. a function called the *cost function* which assigns an m-vector (tuple) $J_1, \dots, J_m$ to each terminal vertex (outcome) of the game $\mathcal{G}$, where J_i denotes the cost of player i, $\mathcal{M} = \{1, \dots, m\}$,
3. each non-terminal vertex of $\mathcal{G}$ is partitioned into $m + 1$ possible sets, $\mathcal{S}_0, \mathcal{S}_1, \dots, \mathcal{S}_m$ called the player sets, where $\mathcal{S}_0$ stands for the choice of chance (nature),
4. each vertex of $\mathcal{S}_0$ has a probability distribution over the edges leading from it,
5. the vertices of each player $\mathcal{S}_i$, $i = 1, \dots, m$ are partitioned into disjoint subsets known as *information sets*, $\mathcal{S}_{i,j}$, such that two vertices in the same information set have the same number of immediate (choices/edges) followers and no vertex can follow another vertex in the same information set.

As a consequence of (5) a player knows which information set he is in but not which vertex of the information set.

A player i is said to have *perfect information* in a game $\mathcal{G}$ is each information set for this player consists of one element. The game $\mathcal{G}$ in extensive form is said to have *perfect information* if every player has perfect information. A *pure strategy* for player i denoted by u_i is defined as a function which assigns to each of player's i information sets $\mathcal{S}_{i,j}$, one of the edges leading from a representative vertex in this set $\mathcal{S}_{i,j}$. As before we denote by Ω_i the set of all pure strategies of player i, $u_i \in \Omega_i$ and by $\mathbf{u} = (u_1, \dots, u_m)$ the m-tuple of all players strategies, with $\mathbf{u} \in \Omega = \Omega_1 \times \cdots \times \Omega_m$. A game in extensive form is *finite* if it has a finite number of vertices, hence each player has only a finite number of strategies. Let us look at a couple of examples.

Example 2.2 In the game of Matching Pennies (see Fig. 2.1) player 1 chooses "heads" (H) or "tails" (T), player 2, not knowing this choice, also chooses between H or T. If the two choose alike (matching) than player 2 wins 1 cent from player 1 (hence $+1$ for player 2 and -1 for player 1); else player 1 wins 1 cent from player 2 (reverse case). The game tree is shown below with vectors at the terminal vertices indicating the cost function, while the number near vertices denote the player to

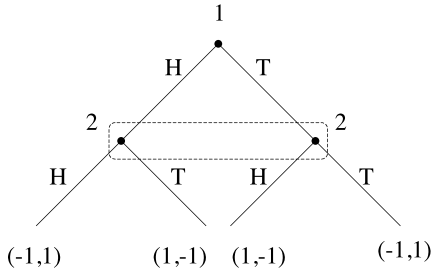

Fig. 2.1 Game of matching pennies in extensive form

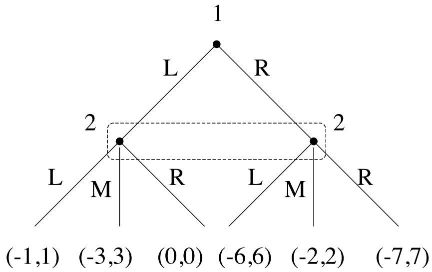

Fig. 2.2 Game of matching pennies in extensive form

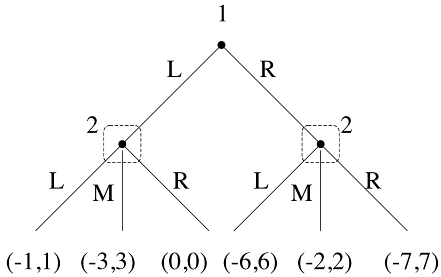

Fig. 2.3 Game of matching pennies in extensive form

whom the move corresponds. The dotted (shaded) area indicates moves in the same information set.

The next two figures show other two zero-sum game examples which differ by the information available to player 2 at the time of its play (information set), denoted by the shaded area (dotted). In the first case, Fig. 2.2, the two possible nodes of player 2 are in the same information set, implying that even though player 1 acts before player 2 does, player 2 does not have access to it s opponent decision. This means that at the time of its play, player 2 does not know at which node (vertex) he is. This is as saying that both players act simultaneously. The extensive form in the second case Fig. 2.3, admits a different matrix game in normal form. In this case each node of player 2 is included in a separate information set, i.e., has perfect information as to which branch of the tree player 1 has chosen.

2.4 Games in Normal Form

Games in normal form (strategic form) model scenarios in which two or more players must make a one-time decision simultaneously. These games are sometimes referred to a one-shot game, simultaneous move games. The *normal form* is a more condensed form of the game, stripped of all features but the choice of each player's pure strategies, and it is more convenient to analyze. The fact that all players make their choice of strategy simultaneously has nothing to do with a temporal constraint, but rather with a constraint on the information structure particular to this type of game. The information structure of a game is a specification of how much each player knows at the time he chooses his strategy. For example, in Stackelberg games [20], where there are leaders and followers, some players (followers) choose their strategies only after the strategic choices made by the leaders have already been revealed.

In order to describe a normal-form game we need to specify players' strategy spaces and cost functions. A strategy space for a player is the set of all strategies available to that player, where a strategy is a complete plan of action for every stage of the game, regardless of whether that stage actually arises in play. A *cost function* of a player is a mapping from the cross-product of players' strategy spaces to that player's set of costs (normally the set of real numbers). We will be mostly concerned with these type of normal-form games herein.

For any strategy profile $\mathbf{u} \in \Omega$, where $\mathbf{u} = (u_1, \dots, u_m) \in \Omega$ is the m-tuple of players' pure strategies and $\Omega = \Omega_1 \times \dots \times \Omega_m$ is the overall *pure strategy space*, let $J_i(\mathbf{u}) \in R$ denote the associated cost for player i that depends on all players' strategies.These costs depend on the context: in economics they represent a firm's profits or a consumer's (von Neumann–Morgenstern) utility, while in biology they represent the fitness (expected number of surviving offspring). All these real numbers $J_i(\mathbf{u})$, $i \in \mathcal{M}$, form the *combined pure strategy* vector cost function of the game, $\mathbf{J} : \Omega \to R_m$. A normal-form game $\mathcal{G}(\mathcal{M}, \Omega_i, J_i)$ is defined by specifying Ω_i and J_i. It is possible to tabulate functions J_i for all possible values of $u_1, \dots, u_m \in \Omega$ either in the form of a relation (easier for continuous or infinite games when Ω is a continuous set), or, as an m-dimensional array (table) in the case of finite games (when Ω is finite set). In this latter case and when $m = 2$ this reduces to a matrix whose size is given by the number of available choices for the two players and whose elements are pairs of real numbers corresponding to outcomes (costs) for the two players. Let us look at a few examples for $m = 2$, where we shall use rows for player 1 as the columns for player 2. Hence entry (j, k) indicates the outcome of player 1 using the j pure strategy and player 2 using k strategy.

Example 2.3 (*Matching Pennies*) Consider the game of Matching Pennies above, where each player has two strategies "Heads" (H) or "Tail" (T). The normal form of this game is described by the matrix

		player 2 H	T
player 1	H	$(-1, 1)$	$(1, -1)$
	T	$(1, -1)$	$(-1, 1)$

or given as the matrix

$$M = \begin{bmatrix} (-1, 1) & (1, -1) \\ (1, -1) & (-1, 1) \end{bmatrix}$$

Most of the times instead of M we shall use a pair of cost matrices (A, B) to indicate the outcome for each player separately, matrix A for player 1 and matrix B for player 2. For the above game this simply means the pair of matrices (A, B) where

$$A = \begin{bmatrix} -1 & 1 \\ 1 & -1 \end{bmatrix}, \qquad B = \begin{bmatrix} 1 & -1 \\ -1 & 1 \end{bmatrix}$$

One can transform any game in extensive form into an equivalent game in normal form, so we shall restrict most of our theoretical development to games in normal form only.

2.5 Game Features

Depending of various features of the game one could classify them in different categories. Below we briefly discuss such classification depending on the competitive nature of the game, the knowledge/information available to the players, and the number of times the game is repeated.

2.5.1 Strategy Space: Matrix vs. Continuous Games

In a matrix game with m players, each player i has a finite number of discrete options to choose from, i.e., there are n_i possible actions, so that the set of its actions is simply identified with a set of indices $M_i = \{1, \dots, n_i\}$ corresponding to these possible actions. Then one considers the action $u_i \in \Omega_i$ with the action sets being defined as $\Omega_i = \{\mathbf{e}_1, \dots, \mathbf{e}_{n_i}\}$, and $\mathbf{e}_j$ being the jth unit vector in $\mathbb{R}^{n_i}$. Given the action $\mathbf{u} \in \Omega$ chosen by all players, player i has a cost matrix $A_i \in \mathbb{R}^{n_1 \times \cdots \times n_m}$, $\forall i \in \mathcal{M}$, that defines his cost by

$$J_i(\mathbf{u}) = [A_i]_{u_1, \dots, u_m} \in \mathbb{R}, \quad i \in \mathcal{M}$$

This is easiest seen in the case of a matrix game when we can explicitly write the cost functions as

$$J_1(\mathbf{u}) = u_1^T A u_2, \qquad J_2(\mathbf{u}) = u_1^T B u_2 \tag{2.1}$$

Such as two-player matrix game is the Matching Pennies game in Example 2.3. A symmetric game is a game where the payoffs for playing a particular strategy depend only on the other strategies employed, not on who is playing them. If the identities of the players can be changed without changing the payoff to the strategies, then a game is symmetric and this corresponds to $B = A^T$. Many of the commonly studied 2×2 games are symmetric. The standard representations of the

Chicken game, Prisoner's Dilemma game, and the Stag Hunt game are all symmetric games [20].

Unlike matrix games, where players have a finite set of actions, in a continuous game each player can choose its action from a continuum of (possibly vector-valued) alternatives, that is, $\Omega_i \subset \mathbb{R}^{n_i}$. We shall review results for both matrix games and continuous games in the next two chapters, but most of the games we shall consider afterwards are continuous games.

2.5.2 Mixed vs. Pure Strategy Games

A strategy can be regarded as a rule for choosing an action, depending on external conditions. Once such a condition is observed, the strategy is implemented as an action. In the case of mixed strategies, this external condition is the outcome of some randomization process. Consider an m-player matrix game and denote by $x_{i,j}$ the probability that player i will choose action j from n_i his available alternatives in Ω_i. Then a mixed strategy x_i is defined as the vector composed of the probabilities associated with available actions, i.e., $x_i = [x_{i,j}]$, $j = 1, \dots, n_i$, $x_i \in \Delta_i$, $i \in \mathcal{M}$ where

$$\Delta_i := \left\{ x_i \in \mathbb{R}^{n_i} \mid x_{i,j} \geq 0, \quad \sum_{j=1}^{n_i} x_{i,j} = 1 \right\}, \quad \forall i \in \mathcal{M}$$

is a simplex. In some cases actions are pure, or independent of any external conditions, and the strategy space coincides with the action space. In discussing games in pure strategies we shall use the term "strategy" and "action" interchangeably to refer to some $\mathbf{u} \in \Omega$, and the game $\mathcal{G}$ can simply be specified as the pair $\mathcal{G}(\mathcal{M}, \Omega_i, J_i)$. This will be the case considered throughout most of the monograph.

2.5.3 Competitive Versus Cooperative

A *cooperative game* is one in which there can be cooperation between the players and/or they have the same cost (also called team games). A *noncooperative game* is one where an element of competition exists and among these we can mention coordination games, constant-sum games, and games of conflicting interests. We give below a few such examples.

2.5.3.1 Coordination Games

In coordination games, what is good for one player is good for all players. An example coordination game in normal form is described by

$$M = \begin{bmatrix} (-3,-3) & (0,0) \\ (0,0) & (-4,-4) \end{bmatrix}$$

In this game, players try to coordinate their actions. The joint action $(j,k) = (2,2)$ is the most desirable (least cost), but the joint action $(j,k) = (1,1)$ also produces negative costs to the players. This particular game is called a pure coordination game since the players always receive the same payoff.

Other coordination games move more toward the domain of games of conflicting interest. For example, consider the Stag Hunt game: stag hare (we shall come back to this example)

$$M = \begin{bmatrix} (-4,-4) & (0,-1) \\ (-1,0) & (-1,-1) \end{bmatrix}$$

In this game, each player can choose to hunt stag (first row or first column) or hare (second row or second column). In order to catch a stag (the biggest animal, hence the bigger payoff or lowest cost of -4), both players must choose to hunt the stag. However, a hunter does not need help to catch a hare, which yields a cost of -1. Thus, in general, it is best for the hunters to coordinate their efforts to hunt stag, but there is considerable risk in doing so (if the other player decides to hunt hare). In this game, the costs (payoffs) are the same for both players when they coordinate their actions, but their costs are not equal when they do no coordinate their actions.

2.5.3.2 Constant-Sum Games

Constant-sum games are games in which the sum of the players' payoffs sum to the same number. These games are games of pure competition of the type "my gain is your loss". *Zero-sum games* are particular example of these games, which in terms of the two-players cost matrices can be described by $B = -A$. An example of such game is the Rock, Paper, and Scissors game with the matrix form

$$M = \begin{bmatrix} (0,0) & (1,-1) & (-1,1) \\ (-1,1) & (0,0) & (1,-1) \\ (1,-1) & (-1,1) & (0,0) \end{bmatrix}$$

2.5.3.3 Games of Conflicting Interests

These fall in between constant-sum games and coordination games and cover a large class, whereby the players have somewhat opposing interests, but all players can benefit from making certain compromises. One can say that people (and learning algorithms) are often tempted to play competitively in these games (both in the real world and in games), though they can often hurt themselves by doing so. However, on the other hand, taking an extreme cooperative approach (same actions) can lead to similarly bad (or worse) payoffs (high costs). One of the most celebrated games of this type is the Prisoners' Dilemma game, with the choices of to "Confess" (cooperate) or "Don't Confess" (defect) as the actions of two prisoners (players) put in separate cells. If they both confess each they each receive 3 years in prison. If only one confesses, he will be freed, used as witness, and the other will be convicted and receive 8 years in prison. If neither confesses they will be convicted of a minor offense and receive each only 1 year. The normal form (strategic form) of this game is described by the matrix

$$M = \begin{bmatrix} (-3,-3) & (0,-8) \\ (-8,0) & (-1,-1) \end{bmatrix}$$

or, as pair of cost matrices (A, B) for the two players,

$$A = \begin{bmatrix} -3 & 0 \\ -8 & -1 \end{bmatrix}, \qquad B = \begin{bmatrix} -3 & -8 \\ 0 & -1 \end{bmatrix}$$

2.5.4 Repetition

Any of the previously mentioned kinds of game can be played any number of times between the same players, and the game can be the same at each play or can be state-dependent.

2.5.4.1 One-Shot Games

In *one-shot games*, players interact for only a single round (or stage). Thus, in these situations there is no possible way for players to reciprocate (by inflicting punishment or rewards) thereafter.

2.5.4.2 Repeated Games

In *repeated games*, players interact with each other for multiple rounds (playing the same game). In such situations, players have opportunities to adapt to each others' behaviors (i.e., "learn") in order to try to become more successful. There can be finite-horizon repeated games where the same game is repeated a fixed number of times by the same players, or infinite-horizon games in which the play is repeated indefinitely.

2.5.4.3 Dynamic Games

The case where the game changes when players interact repeatedly is what can be called a *repeated dynamic game*, characterized by a state. These are also called differential games. Unlike a repeated game where the agents play the same game every time, in a dynamic game the state of the game influences the play and the outcome. Important in this class are the so called *stochastic games*, which are extensions of Markov decision processes to the scenario with m multiple players, where probabilistic transitions are modeled. We shall not cover these types of game in this monograph.

2.5.5 Knowledge Information

Depending on the amount of information a player has different plays and outcomes may be possible. For example, does an player know the costs (or preference orderings) of other players? Does the player know its own cost (payoff) matrix? Can he view the actions and costs of other players? All of these (and other related) questions are important as they can help determine how the player should learn and act. Theoretically, the more information an player has about the game, the better he should be able to do. In short, the information an player has about the game can vary along the following dimensions: knowledge of the player's own actions; knowledge of the player's own costs; knowledge of the existence of other players; knowledge of the other players' actions; knowledge of the other players' costs and in case learning is used, knowledge of the other players' learning algorithms.

In a game with *complete information* each player has knowledge of the payoffs and possible strategies of other players. Thus, *incomplete information* refers to situations in which the payoffs and strategies of other players are not completely known. The term *perfect information* refers to situations in which the actual actions taken by associates are fully observable. Thus, *imperfect information* implies that the exact actions taken by associates are not fully known.

2.6 Solution Concepts

Given a game's specification $\mathcal{G}(\mathcal{M}, \Omega_i, J_i)$ an important issue is to predict how the game will be played, i.e., to determine its outcome. These predictions are called solutions, and describe which strategies will be adopted by players, therefore predicting the result of the game. A solution concept briefly describes how to use a certain set of mathematical rules to decide how to play the game. Various solution concepts have been developed, in trying to indicate/predict how players will behave when they play a generic game. Herein we only introduce these solution concepts in short.

2.6.1 Minimax Solution

One of the most basic properties of every game is the *minimax solution* (or minimax strategy), also called security strategy. The minimax solution is the strategy that minimizes a player's maximum expected loss (cost). There is an alternate set of terminology we can use (often used in the literature as we mentioned before). Rather than speak of minimizing our maximum expected loss, we can talk of maximizing our minimum expected payoff. This is known as the maximin solution. Thus, the terms minimax and maximin can be used interchangeably. The minimax solution is an essential concept for zero-sum games.

Let us look at the Prisoner's Dilemma matrix game above. In the prisoner's dilemma, both players are faced with the choice of cooperating or defecting. If both players cooperate, they both receive a relatively low cost (which is -3 in this case). However, if one of the players cooperates and the other defects, the defector receives a very low cost (-8 in this case) (called the temptation cost), and the cooperator receives a relatively high cost (0 in this case). If both players defect, then both receive a higher cost (which is -1 in this case). So what should you do in this game? Well, there are a lot of ways to look at it, but if you want to play conservatively, you might want to invoke the minimax solution concept, which follows from the following reasoning. If you play cooperate, the worst you can do is get a cost of 0 (thus, we say that the security of cooperating is 0). Likewise, if you play defect, the worst you can do is get a cost of -1 (security of defecting is -1). Alternately, we can form a mixed strategy over the two actions. However, it just so happens in this game that no mixed strategy has higher security than defecting, so the minimax strategy in this game is to defect. This means that the minimax value (which is the maximum cost one can incur when plays the minimax strategy) is -1.

However, even though the minimax value is the lowest cost you can guarantee yourself without the cooperation of your associates, you might be able to do much better on *average* than the minimax strategy if you can either outsmart your associates or get them to cooperate or compromise with you (in a game that is not fully competitive). So we need other solution concepts as well.

2.6.2 Best Response

Another basic solution concept in multi-player games is to play the strategy that gives you the lowest cost given your opponents' strategies. That is exactly what the notion of the *best response* suggests. Suppose that you are player i, and your opponents' play $\mathbf{u}_{-i}$. Then the your best response in terms of pure strategies is u_i^* such that

$$J_i\left(\mathbf{u}_{-i}, u_i^*\right) \leq J_i(\mathbf{u}_{-i}, u_i), \quad \forall u_i \in \Omega_i$$

In the case of mixed strategies, assuming your opponents' play the strategy $\mathbf{x}_{-i}$, your best response is the strategy x_i^* such that

$$J_i\left(\mathbf{x}_{-i}, x_i^*\right) \leq J_i(\mathbf{x}_{-i}, x_i), \quad \forall x_i \in \Delta_i$$

where Δ_i is the probability simplex. The best-response idea has had a huge impact on learning algorithms. If you know what your other players are going to do, why not get the lowest cost (highest payoff) you can get (i.e., why not play a best response)? Taking this one step further, you might reason that if you think you know what other players are going to do, why not play a best response to that belief? While this obviously is not an unreasonable idea, it has two problems. The first problem is that your belief may be wrong, which might expose you to terrible risks. Secondly, this "best-response" approach can be quite unproductive in a repeated game when other players are also learning/adapting [48].

2.6.3 Nash Equilibrium Solution

We now introduce briefly a most celebrated solution concept for a N-player noncooperative game $\mathcal{G}$. John Nash's identification of the *Nash equilibrium* concept has had perhaps the single biggest impact on game theory. Simply put, in a Nash equilibrium, no player has an incentive to unilaterally deviate from its current strategy. Put another way, if each player plays a best response to the strategies of all other players, we have a Nash equilibrium.

We will discuss the extent to which this concept is satisfying by looking at a few examples later on.

Definition 2.4 Given a game $\mathcal{G}$ a strategy N-tuple (profile) $\mathbf{u}^* = (u_1^*, \ldots, u_N^*)$ is said to be a *Nash equilibrium* (or in equilibrium) if and only if

$$J_i\big(u_1^*, \ldots, u_N^*\big) \le J_i\big(u_1^*, \ldots, u_{i-1}^*, u_i, u_{(i+1)}^*, \ldots, u_N^*\big), \quad \forall u_i \in \Omega_i, \ \forall i \in \mathcal{M} \tag{2.2}$$

or, in compact notation,

$$J_i\big(\mathbf{u}_{-i}^*, u_i^*\big) \le J_i\big(\mathbf{u}_{-i}^*, u_i\big), \quad \forall u_i \in \Omega_i, \ \forall i \in \mathcal{M}$$

where $\mathbf{u}^* = (\mathbf{u}_{-i}^*, u_i^*)$ and $\mathbf{u}_{-i}^*$ denotes $\mathbf{u}^*$ of all strategies except the ith one.

Thus $\mathbf{u}^*$ is an equilibrium if no player has a positive incentive for *unilateral* chance of his strategy, i.e., assuming the others keep their same strategies. In particular this means that once all choices of pure strategies have been revealed no player has any cause for regret (hence the point of no regret concept). A similar definition holds for $\mathbf{x}$ mixed strategies (as seen above in the best response).

Example 2.5 Consider the game with normal form

		player 2	
		$u_{2,1}$	$u_{2,2}$
player 1	$u_{1,1}$	$(3, 1)$	$(0, 0)$
	$u_{1,2}$	$(0, 0)$	$(1, 3)$

and note that both $(u_{1,1}, u_{2,1})$ and $(u_{1,2}, u_{2,2})$ are equilibrium pairs. For matrix games we shall use the matrix notation and for the above we will say that $(3, 1)$ and $(1, 3)$ are equilibria.

If we look at another game (a coordination game),

$$M = \begin{bmatrix} (-3, -3) & (0, 0) \\ (0, 0) & (-1, -1) \end{bmatrix}$$

By inspection as in the above, we can conclude that both the joint actions $(j, k) = (1, 1)$ and $(j, k) = (2, 2)$ are Nash equilibria since in both cases, neither player can benefit by unilaterally changing its strategy. Note, however, that this illustrates that not all Nash equilibria are created equally. Some give better costs than others (and some players might have different preference orderings over Nash equilibrium).

While all the Nash equilibria we have identified so far for these two games are *pure strategy* Nash equilibrium, they need not be so. In fact, there is also a third Nash equilibrium in the above coordination game in which both players play mixed-strategies. The next chapter we shall formally review this extension.

Here are a couple more observations about the Nash equilibrium as a solution concept:

- In constant-sum games, the minimax solution is a Nash equilibrium of the game. In fact, it is the unique Nash equilibrium of constant-sum games as long is there is not more than one minimax solution (which occurs only when two strategies have the same security level).
- Since a game can have multiple Nash equilibrium, this concept does not tell us how to play a game (or how we would guess others would play the game). This poses another question: Given multiple Nash equilibria, which one should (or will) be played? This leads to considering refinements of Nash equilibria.

Strategic dominance is another solution concept that can be used in many games. Loosely, an action is strategically dominated if it never produces lower costs (higher payoffs) and (at least) sometimes gives higher costs (lower payoffs) than some other action. An action is strategically dominant if it strategically dominates all other actions. We shall formally define this later on. For example, in the Prisoner's Dilemma (PD) game, the action defect strategically dominates cooperate in the one-shot game. This concept of strategic dominance (or just dominance, as we will sometimes call it) can be used in some games (called iterative dominance solvable games) to compute a Nash equilibrium.

2.6.4 Pareto Optimality

One of the features of a Nash equilibrium (NE) is that in general it does not correspond to a socially optimal outcome. That is, for a given game it is possible for all the players to improve their costs (payoffs) by collectively agreeing to choose a strategy different from the NE. The reason for this is that a posteriori some players may choose to deviate from such a cooperatively agreed-upon strategy in order to improve their payoffs further at the group's expense. A Pareto optimal equilibrium describes a social optimum in the sense that no individual player can improve his payoff (or lower his cost) without making at least one other player worse off. Pareto optimality is not a solution concept, but it can be an important attribute in determining what solution the players should play (or learn to play). Loosely, a Pareto optimal (also called *Pareto efficient*) solution is a solution for which there exists no other solution that gives *every* player in the game a higher payoff (lower cost). A PE solution is formally defined as follows.

Definition 2.6 A solution $\mathbf{u}^*$ is *strictly Pareto dominated* if there exists a joint action $\mathbf{u} \in \Omega$ for which $J_i(\mathbf{u}) < J_i(\mathbf{u}^*)$ for all i, and *weakly Pareto dominated* if there exists a joint action $\mathbf{u} \neq \mathbf{u}^* \in \Omega$ for which $J_i(\mathbf{u}) \leq J_i(\mathbf{u}^*)$ for all i.

Definition 2.7 A solution $\mathbf{u}^*$ is *weakly Pareto efficient (PE)* if it is not strictly Pareto dominated and *strictly Pareto efficient (PE)* if it is not weakly Pareto dominated.

Often, a Nash equilibrium (NE) is not Pareto efficient (optimal). Then one speaks of a loss of efficiency, which is also referred to as the Price of Anarchy. An interesting problem is how to design games with improved Nash efficiency, and pricing or mechanism design is concerned with such issues.

In addition to these solution concepts other important ones include the Stackelberg equilibrium [20], which is relevant in games where the information structure plays an important role, and correlated equilibria [48, 98], which case is relevant in games where the randomization used to translate players' mixed strategies into actions are correlated.

2.7 The Rationality Assumption

Given the number of available solution concepts, NE refinements and the apparent arbitrariness with which they may be applied, why would one expect that in an actual noncooperative game players would choose any particular refined NE? This question turns out to be a valid objection, namely the perfect rationality of all the participating agents. In the literature rationality is often discussed, without being precisely defined. One possible formulation is as follows: a player is rational if it consistently acts to improve its payoff without the possibility of making mistakes, has full knowledge of other players' intentions and the actions available to them, and has an infinite capacity to calculate a priori all possible refinements to $\mathrm{NE}(\mathcal{G})$ in an attempt to find the "best one". If a game involves only rational agents, each of whom believe all other agents to be rational, then theoretical results offer accurate predictions of the game outcomes.

2.8 Learning in Classical Games

Yet another game classification is related to this rationality assumption or the lack of it. In this monograph we will be concerned with rational players and one-shot games. A more realistic modeling scenario involves players that are less than rational and a repeated game play. We review here very briefly the conceptual differences for completeness. The reader is referred to extensive references on this topic such as [48]. We will use the term bounded rationality to describe players that do not necessarily have access to full or accurate information about the game, and who have a limited capacity to perform calculations on the information that is available to them. Instead of immediately playing a perfect move, boundedly rational players adapt their strategy based on the outcomes of previous matches [48, 49, 137]. We can refer to this modeling scenario as a classical game with learning. All solution concepts studied in classical game theory remain important in games with learning.

The important problem to study is not only to classify games for which equilibria exist and have favorable properties such as uniqueness, but also, in conjunction, to classify the strategy update laws that yield convergence to these equilibria under repeated play. In the terminology of [158], a game with learning can be said to have an "inner game" (i.e. the underlying classical game $\mathcal{G}$) and an "outer game" (i.e. the dynamics of the strategy update laws).

One of earliest strategy update laws to be studied is fictitious play (FP) [31]. In FP, each player keeps a running average, known as an empirical frequency of his opponent's actions, and chooses his next move as a best response to this average. The term "fictitious play" comes from the consideration of playing under the unrealistic assumption that the opponent is playing a constant strategy, hence he is a fictitious player. It has been proved that under a number of different sufficient conditions FP is guaranteed to converge to one. One of the earlier cases studied was the class of finite, two-player, zero sum games, for which convergence of FP was proved in [127], followed by other results [75].

A continuous time version of fictitious play can be derived by considering the infinitesimal motions of the empirical frequencies [58]. It can be shown that if both players are updating their strategies so that these are the best response to the other one, we may write

$$\dot{u}_i = -u_i + r_i(\mathbf{u}_{-i})$$

This equation known as the best-response (BR) dynamic, and clearly displays the very convenient feature that $\dot{\mathbf{u}} = 0$ if and only if $\mathbf{u}$ has reached a fixed point of the reaction function $r(\mathbf{u})$. By the characterization of NE given above (see (4.5)), we conclude that the set of equilibria of the best-response dynamic coincides precisely with NE($\mathcal{G}$). Most of these studies are done within the framework of evolutionary game theory [132, 142], while this monograph is concerned only with the setup of classical games.

2.9 Notes

As a preface to the remaining chapters, this chapter provided a brief overview of basic game concepts.

http://www.springer.com/978-0-8176-8321-4

Game Theory for Control of Optical Networks
Pavel, L.
2012, XIII, 261 p. 92 illus., 70 illus. in color., Hardcover
ISBN: 978-0-8176-8321-4
A product of Birkhäuser Basel